

Project news/releases:

None

D0:

Things worth noting:

-- Users complained about d0srv085 with file storage problem at the end of last week.

The LTO2 tape drives are being used extremely heavily (there we also problems with the stores of reconstructed files backing up over the weekend). There is also some evidence there may be problems with the node d0srv085 in particular.

-- There was no minutes published yet for the last week's SAM meeting (Apr. 15) ?

-- Apr. 17: A shifter complained about testCalibClient.py not usable anymore (as the instruction told shifters to try: <http://d0db-prd.fnal.gov/sam/sam-shift-guide.html#calibservers>). I've verified and it seemed to fail in statement "import dcoracle". Can we fix this ? --> And it seems that shifters (including myself) can only ssh as d0db in d0dbsrv5 (for user_prd processes), but NOT to d0db@d0dbsrv4 (for farm_prd processes) by using the root principal. Is this "meant" to be like this ? If you want the shifters to help restarting the calibration servers, one should perhaps copy the .k5login to d0dbsrv4 as well.

Unfortunately, this calibration stuff is basically orphaned, and there is no-one to fix it.

-- Shifter noticed that the mover "D41ELTO" went offline which eventually seemed to be due to a tape stuck in drive

Unresolved issues:

#3542: samg submits failing : Can some expert please take a look at this ? Mike said that the failure rate is now 100% and there is no comment so far.

The disk activity on the machine keeps spiking and making it almost totally unresponsive. Not clear why. [Post meeting: now looks like a hardware issue with the RAID array.]

#3537: "SAM stores backed up?" : seemed to be also related d0srv085. Checking with Enstore to see any new status about d0srv085.

#3534: Files not storing : I checked that Herb's file was eventually stored. I just check with him to see whether we can close this ticket.

#3533: authentication problem at cinvestav-station : Checking with Alberto to see whether his problem has been resolved.

#3520: "D0 failures" --- Robert restarted fcpd processes but still wanted to look into why monitoring didn't alert us. Want to close this ticket ??

CDF:

DCache:

4 cells inactive. Has been true all week.

r-fcdfdata320-1,2,3,4

4 cells in trouble : 2 periods, Sat ~2:30 - Mon 10:40, again Monday 15:30 - 17:30. None at 04:00 Tue.

Enstore Library

CDF-9940B 1/18 Bad Mover: 9940B27

CDF-LT03 Over w/e and again on Monday a large number of pending transfers: stage read (~400), general read (~400) and farm (peaking ~ 200).

AB: Requests well distributed across the dCache pools, and are being processed. All LT03 drives are in use, so the requests cannot be filled any faster than they are. The system will just have to keep chewing through them.

OK now

=====

Tue:

Difference between restore handler and queue info 32, probably due to NO ACCESS tape. Apparently restored in pm (Angela).

Alexei Varganov's job backed up r-fcdfdata146-2. Dealt with it manually

Clxam restarted (AB).

sam v8_1_5 edited to be sam v8_1_6 in fcdffcode3 fixes failures on fcdffdata321 uploads.

Wed

Mover 9940B19 returned to service.

Thur

Bad tapes, fixed later in day:

IAB350

IA6383 (IA3683?)

CDF9940B Bad Movers:

9940B11

9940B27

ReadOnlyPools developed queues, moved temporarily to the TestPools group. (Load on the machines is high).

r-fcdfdata137-1

r-fcdfdata137-3

r-fcdfdata316-1

r-fcdfdata317-1

Fri

New bad mover: 9940B23. 3/18, Bad Tape : IA1480

Sat

CDF-9940B Bad movers:

9940B16

9940B23

9940B27

Sun

Nothing new

Mon

Scheduled outage of server 08:00 - 09:00, no problems.

Process clmeta on b0csl23 lost its heartbeat. Restarted it. (AB)

"The machine fcdldata135, which is a CDF dCache pool server is unpingable. Could someone please investigate?"
Rebooted.

Database response time peaks up to 2.5K Then drops but still ~ 1.5K

From Anil Kumar, email:

EM Alert: Critical: fcdfora4.fnal.gov - CPU Utilization is 100%

Tue 9:00 am.

Known issues now:

4 DCache cells inactive

9940B27 is bad mover

Large number of 9940X movers are idle.

fcdldata135 keeps failing - one of only two machines in b pool group; could move it to a larger pool and replace it with a more reliable machine.

There's a query for historical data volumes by month which is failing with an Oracle internal error which Randolph needs to look into.

Need to copy more runsection data from production to development (for determining prescales). Time consuming because no link between database production and development databases (for good reasons). Note that the development database is being used for production purposes here.

Robert and Angela found a problem with 'availability_status retired' generating the wrong query and never listing anything. This is potentially confusing, but does not affect the functionality of retired files. This is the first time we have tried to retire files for real, so this has not been noticed before. Randolph will investigate.

Minos:

Minos status 2008 Apr 22

General

Enstore -

LTO-4 copies continuing, over 11 TB so far to 15 tapes
encp rates from local disk to LTO-4 are about 15 MB/sec.
This is probably due to our small files (300 MB)

DCache -

SAM -

Upgraded station to v6_0_5_24_srm, using sam_products v4_32

IT 3538

New station would not accept stale projects
(dead, but running pmasters)

04/17/08 15:53:13 minos.SM.REVIVER 26613: 17 projects found

Non-compliant application error detected:

operator->() was used on null pointer or nil object reference

smaster: /fnal/ups/prd/orbacus/Linux-2-4/v3_3_4p1GCC-3-1/include/OB/Template.h:557: T*

OBObjVar<T>::operator->() const [with T = SAMStation_FileConsumer]: Assertion `(int)(ptr_ != 0)'failed.

samadmin purge zombie projects would not remove these,

18 candidate projects found in the database...

Determining which projects are still registered in the NameService...

The following 17 projects are still registered in the NameService and are not eligible for termination:

At some point there was an incompatible change to the smaster/pmaster interface. This upgrade may have gone across this boundary. None of the projects were doing anything, so it wasn't a problem, but does raise the question of why they didn't time out due to lack of activity.

Oracle -

We received excellent warranty, end of service life, and cost information regarding our Oracle hosts and disks from Maureen.

Our plan is to purchase continued support for the existing systems, tentatively till end of life in 2011.

Reducing support to 8to17by7 rather than 24x7,
in line with the rest of our offline support,
saving thousands of dollars on the disk maintenance fees.

OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 2071: quadratic project startup time growth

PLAN -

DB:

Announcement:

- Oracle's Quarterly patch was released week of April 14. Need to schedule this patch on dev/int/prd databases.

D0 :

- Oracle's Quarterly patch on d0ora1~d0ofdev1 April 22, Tuesday 9:30 a.m . Expected Downtime ~2 hours.

- May 06, 2008 Oracle's Quarterly patch and Db patch 5648872 for fix "ORA-error stack (07445[opidsa()+480])" on d0ofint1 and d0ofprd1

Alert(s):

-Owner	Object	Type	Extents Used	Extents Allocated	Extents Left
SAMPRD	FI_PV_PV_FK_I	INDEX	9996	10000	4

CDF :

- CPU Replacement on fcdfora4 April 21, 2008 8:00AM

- Oracle's Quarterly patch was deployed successfully on fcdfora1 on Thursday April 17.

- Need to schedule Oracle's Quarterly patch on cdfopr2 database.
- There are 30 sessions by user 'gdl' accessing SAM. Also executing inefficient queries. User has not responded yet.

Alert(s):

MINOS:

Scheduling Oracle's Quarterly patch on minosdev and minosint On April 24

Scheduling Oracle's Quarterly patch on minosprd for May 15

enstore/dcache report

Scheduled Downtimes:

nothing planned

Operations:

o d0en

- Trouble with LTO2 drives over the weekend turned out to be due to problems with 4 tapes: 1 had been entered upside down after its barcode label was replaced, and 3 blanks that, when loaded, cause the drive to report a media error.
- Trouble with 9940B drives on Sunday resolved by restarting the library manager and media changer. Cause is unknown.
- Barcode reader on Robot 2 in ADIC AML/2 was adjusted, but needs to be replaced when we can schedule a 4 hr downtime.

o cdfe

o cdfca

- Problem with non-responsive doors fixed by restarting admin domain.

o stken

- Bad raid disk on stkensrv3 (backups, pageDcache, monitoring). Disk replaced, but array did not rebuild. Service call placed to vendor.
- Bad raid disk on stkensrv1 (pnfs server). Array rebuilt using a hot spare.

o fndca